

Squared Earth Mover's Distance-based Loss for Training Deep Neural Networks

Le Hou

Computer Science Department
Stony Brook University
lehhou@cs.stonybrook.edu

Chen-Ping Yu

Psychology Department
Harvard University
chenpingyu@fas.harvard.edu

Dimitris Samaras

Computer Science Department
Stony Brook University
samaras@cs.stonybrook.edu

Abstract

Deep neural networks (DNNs) especially convolutional neural networks (CNNs) have achieved state-of-the-art performances in many applications, and they have been shown to be especially powerful in multi-class classification tasks. In deep learning, existing DNNs and CNNs are typically trained with a soft-max cross-entropy loss which only considers the ground-truth class by maximizing the predicted probability of the correct label. This cross-entropy loss ignores the intricate inter-class relationships that exist in the data. In this work, we propose to use the exact squared Earth Mover's Distance (EMD) as a loss function for multi-class classification tasks. The squared EMD loss uses the predicted probabilities of all classes and penalizes the miss-predictions accordingly. In experiments, we evaluate our squared EMD loss in ordered-classes datasets such as age estimation and image aesthetic judgment. We also generalize the squared EMD loss to classification datasets with orderless-classes such as the ImageNet. Our results show that the squared EMD loss allows networks to achieve lower errors than the standard cross-entropy loss, and result in state-of-the-art performances on two age estimation datasets and one image aesthetic judgment dataset.

1. Introduction

Deep neural networks (DNNs) have presently become the preferred method for most machine learning applications, due to their ability to automatically learn the optimal features and classifiers from the inputs in an end-to-end fashion. In addition to superior performance as compared to conventional approaches, another reason for the popularity is their wide-range of applicability which includes convolutional neural networks (CNNs) for computer vision [15, 33], recurrent neural networks (RNNs) for natural language processing [34, 22], hybrid networks that combine CNN and RNN layers for speech recognition and audio processing [14, 6], and more. In general, DNNs are trained under one of two general tasks: regression and classification. In a re-

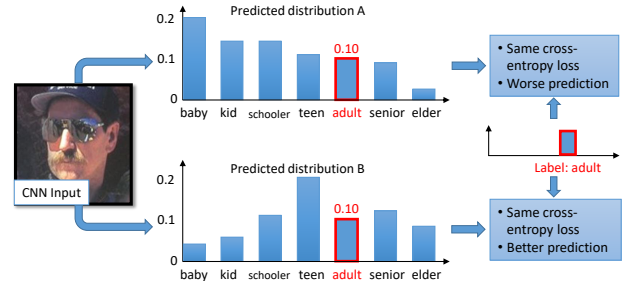


Figure 1. In many classification tasks, there are relationships or even orders between classes. However the cross-entropy loss ignores these relationships and only focus on the predicted probability of the ground truth class. In this example, the two given predicted distributions have the same cross-entropy loss. But clearly the predicted distribution B is preferred.

gression task, the network learns to generate a real-valued output that matches the ground-truth [1, 7]. In a classification task, the network learns to categorize an input to one of the training classes [9, 38, 15, 33, 21].

Multi-class classification is a typical problem for applying a categorization DNN. In the case of a regression problem, transforming the real-valued labels into ordered classes and apply a multi-class classification DNN is also popular, because DNNs are easier to train under a classification framework and often converge to lower errors [13]. Examples of such datasets with ordered-classes include age estimation [9, 10, 12], image aesthetics [16], facial attractiveness prediction [35] and others [25]. Previous state-of-the-art performance on these datasets has been attained via training a multi-class classification with a DNN [29, 25, 8].

To train a multi-class single-label classification network, softmax cross-entropy loss is by far the most popular loss function for the training regime, where the ground-truth is a binary vector consists of a value 1 at the correct class index, and 0s everywhere else [18, 17]. During training, the objective is to minimize the negative log-likelihood of the loss by multiplying the network's predictions to the binary ground-truth vectors. While the softmax cross-entropy has demon-

strated empirical success across all applicable fields, there is a major aspect that the loss function does not take into account: the similarities between the classes. For example (Fig. 1), we want to estimate human age stages from face images. Given an image of an adult, if network A’s output probability of the adult class is 0.1 with a max probability of 0.2 at the baby class, and network B’s output probability of the adult class is 0.1 with a max probability of 0.2 at the teen class, they would receive the same amount of loss from softmax cross-entropy, while network B’s output probabilities are clearly closer to the ground-truth. This is especially true for the ordered-classes datasets, but also can be extended to datasets with orderless-classes.

In this work, we show how the exact squared Earth Mover’s Distance (EMD) [30] can be applied as a superior loss function compared to the standard softmax cross-entropy loss for multi-class classification problems using DNNs. The EMD is also known as the Wasserstein distance [2], which is the minimal cost required to transform one distribution to another [30]. A recent work formulated an approximated Wasserstein loss for supervised multi-class multi-label learning using a linear model, and applied to classification problems with predefined inter-class similarity metrics [11]. In contrast, we focus on multi-class single-label learning using DNNs. We show how an *exact* squared EMD (EMD²) can be used as a loss function for training single-label deep learning models directly on datasets with ordered-classes. We choose to use EMD² instead of EMD as the loss function because it usually converges faster with gradient descent [32, 24]. CNNs with our EMD² loss result in better performance than CNNs with the standard softmax cross-entropy loss, and achieve state-of-the-art results on multiple datasets. Moreover, we extend this approach as a regularization term for ImageNet [31] that has no class ranking relationships, using a self-guided method that allows the DNN to estimate the similarities between classes using its own features.

We make four major contributions in this paper:

1. We show how an exact EMD² loss function can be used to train DNNs in ordered-classes scenarios.
2. We propose a self-guided end-to-end training method with EMD² as regularization to cross-entropy for general case classification using DNNs. This self-guided approach estimates the similarities between classes using the DNN’s own features.
3. To the best of our knowledge, our framework is the first work that demonstrates the full use of EMD² as a loss function in deep learning, for both ordered and general-case multi-class classification.
4. Our proposed method allows networks to achieve state-of-the-art performances on two age-estimation

and one aesthetic judgement datasets when no external training data is used, and an improved performance over the softmax cross-entropy on the general-case extension using ImageNet.

The rest of this paper is organized as follows. Sec. 2 formulates the EMD² loss for classification problems with ordered classes. Sec. 3 shows our method of estimating inter-class relationships between orderless classes, and generalizes the EMD² loss on classification problems with orderless-classes. Sec. 4 shows the advantage of our method experimentally. Sec. 5 concludes this paper.

2. EMD for Ordered-Classes Classification

In this section, we first introduce the standard softmax cross-entropy loss and discuss its drawbacks in detail. Then, we formulate the EMD and show it can be used as a loss function for classification problems with ordered-classes.

2.1. Softmax Cross-Entropy Loss

The softmax cross-entropy loss is a log-likelihood based loss function, that combines a softmax layer with a cross-entropy loss function. For a single-label classification problem with C classes, a network’s softmax layer outputs a probability distribution \mathbf{p} of length C , with its i -th entry \mathbf{p}_i being the predicted probability of the i -th class. The softmax guarantees that $\sum_i \mathbf{p}_i = 1$. We denote the ground truth as a binary vector \mathbf{t} of length C . Also $\sum_i \mathbf{t}_i = 1$. Given a training example, the cross-entropy loss between prediction \mathbf{p} and ground truth vector \mathbf{t} is defined as:

$$E_X(\mathbf{p}, \mathbf{t}) = - \sum_{i=1}^C (\mathbf{t}_i \log(\mathbf{p}_i)). \quad (1)$$

We assume that the k -th class is the ground truth label for single-label classification problems. In other words $\mathbf{t}_k = 1$ and $\mathbf{t}_i = 0$ for $i \neq k$. From Eq. 1, the cross-entropy loss $E_X(\mathbf{p}, \mathbf{t}) = -\log(\mathbf{p}_k)$. Its differentiation is:

$$E'_X(\mathbf{p}, \mathbf{t}) = -\mathbf{p}'_k / \mathbf{p}_k. \quad (2)$$

It is obvious to see that the backpropagation of a DNN with cross-entropy loss only depends on \mathbf{p}_k . We argue that this is less robust compared to a loss function that depends on all entries of \mathbf{p} .

2.2. EMD² Loss

In this section, we first define the problem of ordered-classes classification. We then introduce the Earth Mover’s Distance (EMD) as a similarity measurement between distributions, and explain how an EMD²-based loss function models inter-class similarities. We then show that the exact EMD² can be computed by a closed-form equation in ordered-classes classification problems without any approximation.

2.2.1 Ordered-classes classification

We define ordered-classes as classes that can be represented as real numbers. For example, the human age classes and the preference levels can be described by integers. These classes can be ranked based on their inherent ordering. The difference between ordered-classes classification and regression is that in the problem of ordered-classes classification, the ground truth labels and predictions are discrete. Although ordered-classes can be learned by a regression model, it has been shown in practice that better performance can be achieved using a multi-class classification model [13].

Without loss of generality, we assume that in all ordered-classes classification problems, the classes have already been ranked as $t_1, t_2, t_3, \dots, t_C$. We further assume that for a pair of classes, if they are close to each other (have similar rank positions), then their member images share similar visual features; if they are far away from each other (have very different rank positions), then their member images share little or no visual features, and that these visual features are critical for correctly classifying the images.

2.2.2 Earth Mover's Distance

We assume that a well performing CNN should predict class distributions such that classes closer to the ground truth class should have higher predicted probabilities than classes that are further away. We formulate this using the Earth Mover's Distance (EMD). The EMD is defined as the minimum cost to transport the mass of one distribution (histogram) to the other, using ground distances between classes.

Mass transportation defines the problem of transporting mass from a set of supplier clusters to a set of consumer clusters. Its formal definition [30] is: Let $\mathbf{p} = \{(\mathbf{a}_1, \mathbf{p}_1), (\mathbf{a}_2, \mathbf{p}_2), \dots, (\mathbf{a}_C, \mathbf{p}_C)\}$ be the supplier signature (distribution or histogram) with C clusters (bins), where B_i represents each cluster and \mathbf{p}_i is the mass (value) in each cluster. Let $\mathbf{t} = \{(\mathbf{b}_1, \mathbf{t}_1), (\mathbf{b}_2, \mathbf{t}_2), \dots, (\mathbf{b}_{C'}, \mathbf{t}_{C'})\}$ be the consumer signature. Let \mathbf{D} be the ground distance matrix where its i, j -th entry $D_{i,j}$ is the distance between \mathbf{a}_i and \mathbf{b}_j . Matrix \mathbf{D} is usually defined as the l -norm distance between clusters:

$$D_{i,j} = \|\mathbf{a}_i - \mathbf{b}_j\|_l \quad (3)$$

Let \mathbf{F} be the transportation matrix where its i, j -th entry transports $F_{i,j}$ indicates the mass transported from \mathbf{a}_i to \mathbf{b}_j . A valid transportation satisfies four constraints. First, the amount of mass transported must be positive:

$$F_{i,j} \geq 0 \quad \text{for all } i, j. \quad (4)$$

Second, the amount of mass transported from a supplier cluster \mathbf{p}_i must not exceed its total mass:

$$\sum_{j=1}^{C'} F_{i,j} \leq \mathbf{p}_i \quad \text{for all } i. \quad (5)$$

Third, the amount of mass transported to a consumer cluster \mathbf{t}_j must not exceed its total mass:

$$\sum_{i=1}^C F_{i,j} \leq \mathbf{t}_j \quad \text{for all } j. \quad (6)$$

Finally, the total flow $\sum_{i=1}^C \sum_{j=1}^{C'} F_{i,j}$ must not exceed the total mass that can be transported:

$$\sum_{i=1}^C \sum_{j=1}^{C'} F_{i,j} = \min \left(\sum_{i=1}^C \mathbf{p}_i, \sum_{j=1}^{C'} \mathbf{t}_j \right). \quad (7)$$

Under the constraints defined above, the overall cost of flow \mathbf{F} is defined as:

$$W(\mathbf{b}, \mathbf{t}, \mathbf{F}) = \sum_{i=1}^C \sum_{j=1}^{C'} D_{i,j} F_{i,j}. \quad (8)$$

The EMD between two vectors, denoted as $\text{EMD}(\mathbf{p}, \mathbf{t})$ is the minimum cost of work that satisfies the constraints in Eq. 4, 5, 6, 7, normalized by the total flow.

$$\text{EMD}(\mathbf{p}, \mathbf{t}) = \frac{\sum_{i=1}^C \sum_{j=1}^{C'} D_{i,j} F_{i,j}}{\sum_{i=1}^C \sum_{j=1}^{C'} F_{i,j}}. \quad (9)$$

As described, the EMD computes the similarity between two distributions \mathbf{p} and \mathbf{t} [30]. The EMD differs from other distribution similarity metrics such as KL-divergence [37], in that the EMD utilizes the ground distance matrix \mathbf{D} to define the distances between the clusters (bins) of a distribution.

2.2.3 EMD² loss for ordered-classes classification

EMD has been shown to be equivalent to the Mallows distance which has a closed-form solution [20], if the ground distance matrix \mathbf{D} and distributions \mathbf{p} and \mathbf{t} satisfy certain conditions, as shown in [20]. In the following, we show that the required conditions are satisfied under the ordered-classes classification problems.

The first condition is that the two distributions \mathbf{p} and \mathbf{t} to be compared must have equal mass:

$$\sum_{i=1}^C \mathbf{p}_i = \sum_{j=1}^{C'} \mathbf{t}_j \quad (10)$$

Note that this condition is always satisfied if \mathbf{p} is produced by a softmax layer, as the output vector of a softmax layer is a normalized probability density function that sums to 1. And since the number of classes in the predicted distribution is the same as the target distribution, then $C = C'$.

The second condition is that the ground distance matrix \mathbf{D} must have an one-dimensional embedding. Assuming $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_C$ and $\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_{C'}$ are sorted according to their inherent rank values without loss of generality, this condition can be expressed as:

$$\mathbf{D}_{i,j} = S(j - i), \quad (11)$$

for a constant S and all i, j that $i \leq j$. Clearly, this assumption can always be satisfied in an ordered-class classification problems.

The third and final condition is that the distributions to be compared must be sorted vectors. This condition is always satisfied since we assumed $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_C$ and $\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_{C'}$ are sorted without loss of generality. Then, based on the conclusion by Levina *et al.* [20], the normalized EMD can be computed exactly and in closed-form with ordered-classes prediction and target distributions:

$$\text{EMD}(\mathbf{p}, \mathbf{t}) = \left(\frac{1}{C}\right)^{\frac{1}{l}} \|\text{CDF}(\mathbf{p}) - \text{CDF}(\mathbf{t})\|_l, \quad (12)$$

where $\text{CDF}(\cdot)$ is a function that returns the cumulative density function of its input.

We use $l = 2$ for Euclidean distance and also for \mathbf{D} in Eq. 3. Dropping the normalization term, we obtain the final EMD^2 loss E_E as:

$$E_E(\mathbf{p}, \mathbf{t}) = \sum_{i=1}^C \left(\text{CDF}_i(\mathbf{p}) - \text{CDF}_i(\mathbf{t}) \right)^2, \quad (13)$$

where $\text{CDF}_i(\mathbf{p})$ is the i -th element of the CDF of \mathbf{p} . This equation is directly applicable to ordered-classes classification problems. Note that we choose to use EMD^2 instead of EMD as the loss function because it usually converges faster and is easier to optimize with gradient descent [32, 24].

2.2.4 Derivative of EMD^2 loss

To derive the derivatives of E_E with respect to the network parameters, we first rewrite E_E as:

$$\begin{aligned} E_E(\mathbf{p}, \mathbf{t}) = & (\mathbf{p}_1 - \mathbf{t}_1)^2 + (\mathbf{p}_1 + \mathbf{p}_2 - \mathbf{t}_1 - \mathbf{t}_2)^2 + \dots \\ & + \left(\sum_{i=1}^C \mathbf{p}_i - \sum_{j=1}^C \mathbf{t}_j \right)^2. \end{aligned} \quad (14)$$

Each prediction \mathbf{p}_i is a function of the network parameters. Differentiating $E_E(\mathbf{p}, \mathbf{t})$ with respect to the network param-

eters yields:

$$\begin{aligned} E'_E(\mathbf{p}, \mathbf{t}) = & 2 \sum_{n=1}^C \left(\left(\sum_{i=1}^n \mathbf{p}_i - \sum_{j=1}^n \mathbf{t}_j \right) \left(\sum_{i=1}^n \mathbf{p}'_i \right) \right) \\ = & 2\mathbf{p}'_1 \left(\sum_{i=1}^C (C - i + 1) (\mathbf{p}_i - \mathbf{t}_i) \right) \\ & + 2\mathbf{p}'_2 \left(\sum_{i=1}^C (C - i + 1) (\mathbf{p}_i - \mathbf{t}_i) - \mathbf{p}_1 + \mathbf{t}_1 \right) + \dots \\ & + 2\mathbf{p}'_C \left(\sum_{i=1}^C (C - i + 1) (\mathbf{p}_i - \mathbf{t}_i) - \sum_{i=1}^{C-1} (C - i) (\mathbf{p}_i - \mathbf{t}_i) \right) \end{aligned} \quad (15)$$

The coefficients of \mathbf{p}'_i can be propagated using the standard backpropagation method.

Comparing Eq. 15 with Eq. 2, we can see that the backpropagation of a network trained with cross-entropy loss is only based on \mathbf{p}_k and \mathbf{p}'_k the ground truth element, whereas the backpropagation of a network trained with EMD^2 loss is based on all elements of \mathbf{p} and \mathbf{p}' .

3. EMD for Orderless-Class Classification

Sec. 2.2.3 showed the formulation of EMD^2 loss for ordered classes, based on the properties of the ground distance matrix \mathbf{D} inherited from the natural ordering of the classes. In the case of orderless-classes datasets, the relationship between the classes lays in an unknown high dimensional space. In other words, the matrix \mathbf{D} for these problems is unknown. In this section, we show how to compute a ground distance \mathbf{D} using empirical evidence. Furthermore, we propose a method that computes the EMD^2 -based loss between prediction \mathbf{p} and ground truth \mathbf{t} with $O(C)$ time complexity for single-label classification problems. Finally, for the orderless-classes classification problems, we propose to use EMD^2 as a regularization term to the cross-entropy loss.

3.1. Estimating ground distances

We focus on estimating the ground distance matrix \mathbf{D} expressed in Eq. 3. Note that for classification problems, the predicted classes (supplier signatures) and the ground truth classes (consumer signatures) are the same set of classes. In other words $\mathbf{a}_i := \mathbf{b}_i$ for all i . We will use \mathbf{a} to indicate the same set of classes in the future. A straightforward method of estimating \mathbf{D} is to estimate all \mathbf{a}_i directly and compute \mathbf{D} directly from Eq. 3. To estimate each \mathbf{a}_i , one can extract features on all instances of the i -th class, and compute the centroid of all feature vectors.

Since a DNN or a CNN in most multi-class classification cases, automatically computes feature vectors of each instance, we use these CNN features directly to compute

the location of each class in the high dimensional space. To estimate the representation of the i -th class, we take the CNN that is being trained, and extract its own second-to-last layer neural responses from all instances that belong to the i -th class, and use them as the features for the class members. We then $L1$ normalize these CNN features. Finally, we perform element-wise averaging on all extracted feature vectors. Because CNNs learn to *linearly* separate classes with the second-to-last layer features (there is no subsequent non-linearity), it is sufficient to simply *average* the feature vectors to estimate the location of the class in the high dimensional space. This approach allows the network to guide itself over the course of training, without relying on features extracted using a separate pre-trained network.

However, we observe in practice that the CNN features cannot provide sufficient class separation unless the network itself has already partially converged. As a result, the high dimensional locations of these classes are very close to each other. In other words, the estimated matrix $\bar{\mathbf{D}}$ only contains a few different values. To address this, we map each row of $\bar{\mathbf{D}}$ onto a set of uniformly distributed values. Denoting the transformed matrix as \mathbf{B} , this operation is formulated as:

$$\mathbf{B}_{i,j} = \frac{1}{C} \mathbf{R}(\bar{\mathbf{D}}_{i,j}, \{\bar{\mathbf{D}}_{i,1}, \bar{\mathbf{D}}_{i,2}, \dots, \bar{\mathbf{D}}_{i,C}\}), \quad (16)$$

where $\mathbf{R}(\bar{\mathbf{D}}_{i,j}, \{\bar{\mathbf{D}}_{i,1}, \bar{\mathbf{D}}_{i,2}, \dots, \bar{\mathbf{D}}_{i,C}\})$ returns the number of elements in the set $\{\bar{\mathbf{D}}_{i,1}, \bar{\mathbf{D}}_{i,2}, \dots, \bar{\mathbf{D}}_{i,C}\}$ that is smaller than $\bar{\mathbf{D}}_{i,j}$. After this transformation, all entries of \mathbf{B} are between 0 and 1, and our final estimation of the ground distance matrix \mathbf{D} is a symmetric matrix obtained below:

$$\mathbf{D} = (\mathbf{B} + \mathbf{B}^T)/2. \quad (17)$$

Note that based on this definition, $\mathbf{D}_{i,i} = 0$ for all i .

3.2. Self-guided EMD² regularization

We now focus on calculating the EMD between \mathbf{p} and \mathbf{t} , defined by Eq. 9. In the case of single-label classification, the consumer's mass vector (target distribution) \mathbf{t} is a binary vector with only the index of the ground truth class k equals to 1: $\mathbf{t}_k = 1$. According to the constraint defined by Eq. 6, all mass must be transferred to the k -th cluster. In other words, the transportation matrix must satisfy $\mathbf{F}_{i,j} = 0$ if $j \neq k$, otherwise $\mathbf{F}_{i,j} = \mathbf{p}_i$. The resulting EMD is expressed below:

$$\text{EMD}(\mathbf{p}, \mathbf{t}) = \frac{\sum_{i=1}^C \mathbf{p}_i \mathbf{D}_{i,k}}{\sum_{i=1}^C \mathbf{p}_i} = \sum_{i=1}^C \mathbf{p}_i \mathbf{D}_{i,k}. \quad (18)$$

The time complexity computing EMD by Eq. 18 is $O(C)$.

The exact EMD defined by Eq. 18 can be directly used as a loss function. Its derivative with respect to network parameters is:

$$\text{EMD}'(\mathbf{p}, \mathbf{t}) = \sum_{i=1}^C \mathbf{p}_i' \mathbf{D}_{i,k}. \quad (19)$$

In practice the optimization does not converge to a desired local optimum using Eq. 18 directly. We observed that using Eq. 19 for gradient descent ends up lowering the predicted probabilities of all classes except the ground truth class, which leads the model to converge to a local minimum with uniformly distributed predictions. To address this optimization problem, we modify Eq. 18 and use it as a regularizer instead of a stand-alone loss function. Additionally, for faster optimization with gradient descent, we use \mathbf{p}_i^2 instead of \mathbf{p}_i as the mass for each supplier cluster. Our hybrid loss with EMD²-based regularization for the orderless-classes classification problem is then defined as:

$$\mathbf{E}_H(\mathbf{p}, \mathbf{t}) = \mathbf{E}_X(\mathbf{p}, \mathbf{t}) + \lambda \sum_{i=1}^C \mathbf{p}_i^2 (\mathbf{D}_{i,k}^\omega + \mu), \quad (20)$$

where λ , ω and μ are predefined parameters, such that λ defines the weight of EMD² regularization, the power term ω determines the sensitivity of the ground distance: a very high ω means that EMD only penalizes predictions on classes that are far away from the ground truth class, and μ is the ground distance bias. In our experiments, we use a negative μ so that $\mathbf{D}_{i,k}^\omega + \mu$ is negative, which means that the network will be rewarded for predictions that are closer to the ground truth class.

4. Experiments and Results

In this section, we show implementation details and experimental results. First, we apply our EMD²-based losses on three datasets with ordered-classes: the Adience age estimation dataset [9], the Images of Groups age estimation dataset [12], and the image aesthetics with attributes database (AADB) [16]. We achieve state-of-the-art results on all three datasets when comparing to methods that do not use external training data, outperforming both regression CNNs, and classification CNNs with cross-entropy loss. Additionally, we show the performance of using the self-guided EMD² regularization on ImageNet ([31]) as the extension to orderless-classes classification.

4.1. Implementation Details

We tested the EMD²-based losses on different network architectures including the AlexNet [17], VGG 16-layer network [33], and wide residual network [38]. For optimization, we used stochastic

gradient descent with momentum 0.98 in all experiments. The learning rates were selected from $\{10^{-1.5}, 10^{-2}, 10^{-2.5}, 10^{-3}, 10^{-3.5}, 10^{-4}, 10^{-4.5}\}$ individually for each method on each dataset. We noticed that when using EMD^2 as a regularizer, predicted probability for some classes can be very close to zero, resulting in errors when computing the logarithm of the prediction vector. To solve this, we simply added 1^{-6} to the predicted probabilities of all classes. For experiments on ImageNet, we used the same data augmentation method used by AlexNet [17]. For experiments on all other datasets, we used the following data augmentation methods: during training, we first cropped smaller images from the original images (translation augmentation); second, we perturbed the images' RGB colors slightly; third, the images were randomly flipped horizontally; fourth, we rotated the images by $(-20, 20)$ degrees; finally, we adjusted the aspect ratio by $\pm 10\%$. During testing, we used the average prediction from the center crop and its mirrored image. We used Theano [36] for network implementation. The details of the naming conventions for the individual experiments are as follows:

ALX / VGG / RES The original AlexNet (ALX) [17] by [19], VGG 16-layer network (VGG) [33], and wide residual network (RES) [38] with softmax cross-entropy loss. In this experiment, we trained AlexNets in all experiments from scratch. For the VGG network, we use an ImageNet [33] pretrained network and fine-tuned it on each dataset. For the wide residual network, we used a 40-layer network with identity mapping and bottleneck design that is the same as [38]. We also trained wide residual networks from scratch. The pre-training of VGG was based on the comparison to [29], as the authors in [29] reported their results using a VGG that was pre-trained on ImageNet.

ALX-REG / VGG-REG / RES-REG A regression version of AlexNet (ALX-REG), VGG (VGG-REG), and ResNet (RES-REG). The regression networks are implemented by attaching a fully connected layer with one output neuron with linear activation function, following the conventional regression CNN approach [1].

ALX-EMD / VGG-EMD / RES-EMD Using EMD^2 loss defined by Eq. 13 on ordered-classes classification problems, with AlexNet (ALX-EMD), VGG (VGG-EMD), and ResNet (RES-EMD).

ALX-XE-EMD1 / VGG-XE-EMD1 / RES-XE-EMD1 Testing self-guided EMD^2 regularization training defined by Eq. 20 with AlexNet (ALX-XE-EMD1), VGG (VGG-XE-EMD1), and ResNet (RES-XE-EMD1). For the predefined parameters in Eq. 20, we chose $\omega = 1$ and $\mu = 0.5$. We "jump-started" the

Method	AEM	AEO
Dropout-SVM [9]	45.1	79.5
ALXs-Levi [19]	50.7	84.7
Cascade CNN [3]	52.9	88.5
VGG-DEX [29]	55.6	89.7
VGG-DEX + IMDB-WIKI [29]	64.0	96.6
ALXs	53.0	85.7
ALXs-REG	49.8	85.1
ALXs-EMD	57.0	90.4
ALXs-XE-EMD1	54.2	88.0
ALXs-XE-EMD2	54.8	87.5
VGG	60.9	92.8
VGG-REG	56.5	94.0
VGG-EMD	59.2	92.6
VGG-XE-EMD1	60.7	93.7
VGG-XE-EMD2	61.1	94.0

Table 1. Experimental results on the Adience dataset [9]. The pre-trained VGG network fine-tuned using our proposed loss function (VGG-XE-EMD2) outperformed the state-of-the-art Deep Expectation (VGG-DEX) method [29] on the same dataset. The VGG-DEX + IMDB-WIKI [29] method achieved better results using an external age estimation dataset IMDB-WIKI for training, which is 10 times larger than Adience. The results also show that the EMD^2 -based loss helped the models to converge to better accuracies than both cross-entropy and the regression models.

networks by training with softmax cross-entropy for the first 4 epochs with $\lambda = 0$, as a way to avoid using inaccurately estimated ground distance matrix \mathbf{D} for computing class representations. After 4 epochs we chose a λ such that the EMD^2 term is 3 to 4 times smaller than the cross-entropy term. We found that in ordered-classes datasets, the performance was not sensitive when we changed λ .

ALX-XE-EMD2 / VGG-XE-EMD2 / RES-XE-EMD2

Testing self-guided EMD^2 regularization training defined by Eq. 20, but with $\omega = 2$ and $\mu = 0.25$ and keep other parameters unchanged, using AlexNet (ALX-XE-EMD2), VGG (VGG-XE-EMD2), and ResNet (RES-XE-EMD2).

For experiments on the Adience dataset, following the baseline method [19], we tested a smaller version of AlexNet. We name it as **ALXs**.

4.2. Age Estimation on Adience Dataset

Age estimation using human face images is important for analyzing and understanding human faces [28], and has applications such as face identification [27, 26]. We tested our method on the Adience dataset [9]. The Adience dataset contains 26,000 images in 8 age groups, and a five-fold cross-validation evaluation scheme. For comparison, we

Age Groups	0-2	4-6	8-13	15-20	25-32	38-43	48-53	60-
0-2	0	0.32	0.76	0.79	0.83	0.88	0.94	1.00
4-6	0.32	0	0.32	0.57	0.76	0.81	0.88	0.94
8-13	0.76	0.32	0	0.32	0.50	0.69	0.75	0.81
15-20	0.79	0.57	0.32	0	0.25	0.57	0.69	0.76
25-32	0.83	0.76	0.50	0.25	0	0.38	0.57	0.64
38-43	0.88	0.81	0.69	0.57	0.38	0	0.32	0.44
48-53	0.94	0.88	0.75	0.69	0.57	0.32	0	0.25
60-	1.00	0.94	0.81	0.76	0.64	0.44	0.25	0

Table 2. The learned ground distance matrix \mathbf{D} (Eq. 17) by VGG-XE-EMD2 on the Adience dataset. The age groups that are further away from each other always have larger ground distances than age groups that are closer to each other. The interesting semantic differences were also captured during training (e.g. people in their 15-20s are more similar to 25-32s than to 8-13s).




Figure 2. The AEM and AEO curves on the Adience dataset (best viewed in color). X-axis: number of training epochs. Y-axis: averaged AEM and AEO results by five-fold cross-validation. Our proposed EMD²-based loss functions outperformed the cross-entropy based loss functions significantly.

used a smaller version of AlexNet [19] and fine-tuned a pre-trained VGG 16-layer network [29] as described in the earlier experimental details. We used the conventional accuracy of exact match (AEM%) and with-in-one-category-off match (AEO%) as evaluation metrics.

The results are shown in Tab. 1. The pre-trained VGG network fine-tuned using our proposed EMD²-based loss function (VGG-XE-EMD2) outperformed the state-of-the-art Deep Expectation (VGG-DEX) method [29] on the same dataset. The VGG-DEX + IMDB-WIKI [29] method achieved better results using an external age estimation dataset IMDB-WIKI which is 10 times larger than Adience. Therefore, our method improved the state-of-the-art when training without external face datasets. We show the

AEM and AEO results with respect to the number of training epochs in Fig. 2, and the learned ground distance matrix \mathbf{D} (Eq. 17) by our best performing method VGG-XE-EMD2 on the Adience dataset in Tab. 2. We can see that age groups that are further away from each other always have larger ground distances than age groups that are closer to each other. In addition, the ground distance matrix has also learned asymmetries in similarities between different age groups, e.g., people in their 15-20s are more similar to 25-32s than to 8-13s.

4.3. Age Estimation on Images of Groups Dataset

We further test our method on the Images of Groups dataset [12]. This dataset contains 3,500 training face images and 1,000 testing face images in 7 age groups. We trained wide residual networks [38] from scratch and fine-tuned a pre-trained VGG 16-layer network [29] on this dataset. The results are shown in Tab. 3. Our EMD²-based losses outperformed the cross-entropy loss and the $L2$ loss (regression) in a majority of the cases, with our results achieving a new state-of-the-art on this dataset.

4.4. Image Aesthetics

Assessing image aesthetics automatically has a wide range of applications including automatic photo management [23] and image retrieval [4, 5]. We tested our method on the Image Aesthetics with Attributes Database (AADB) [16]. The AADB dataset contains 8,458 training and 1,000 testing images, labeled as real numbers in $[0.0, 1.0]$ according to the viewers' aesthetic judgments. To transform this dataset into a classification dataset, we discretized the real number labels to 10 aesthetic bins, balancing the number of training images in each bin. During testing, we computed the expected aesthetic scores according to the predicted distributions of 10 aesthetic bins. This gave us real-numbered predictions. We used Spearmans' rank correlation ρ as the evaluation metric, following [16]. We trained wide residual networks [38] from scratch, and fine-tuned a pre-trained VGG 16-layer network [29] on this dataset, as with earlier

Method	AEM	AEO
single-CNN [8]	54	90
multi-CNN [8]	56	92
RES	60.0	91.5
RES-REG	52.8	92.2
RES-EMD	59.3	92.5
RES-XE-EMD1	58.2	91.7
RES-XE-EMD2	60.1	93.1
VGG	64.3	95.6
VGG-REG	60.2	96.6
VGG-EMD	65.0	96.1
VGG-XE-EMD1	63.8	95.4
VGG-XE-EMD2	64.6	96.1

Table 3. Experimental results on the Images of Groups dataset [9]. Our EMD²-based losses outperformed the cross-entropy loss and the l2 loss (regression) based methods in majority of the cases. We achieved new state-of-the-art on this dataset.

Method	Spearman's ρ
ALX-Kong [16]	0.5923
ALX-REG-Kong [16]	0.6239
Reg+Rank+Att+Cont* [16]	0.6782
RES	0.5003
RES-REG	0.5235
RES-EMD	0.5448
RES-XE-EMD1	0.5370
RES-XE-EMD2	0.5147
VGG	0.6283
VGG-REG	0.6096
VGG-EMD	0.6682
VGG-XE-EMD1	0.6371
VGG-XE-EMD2	0.6297
VGG-EMD \times 8	0.6889

Table 4. Experimental results on the image aesthetics with attributes database (AADB) [16]. *: use additional 11 labels of image attributes such as color harmony, and image content information. Our EMD²-based losses outperformed cross-entropy loss and l2 loss (regression) based methods significantly. By averaging the results of eight VGG-EMD networks, the results were better than the various attribute-augmented model [16] with just image data, and achieved a new state-of-the-art result.

Method	Top-1 Error	Top-5 Error
ALX	0.435	0.207
ALX-XE-EMD1	0.422	0.202

Table 5. Experimental results on the ImageNet ILSVRC 2012 dataset [31]. The AlexNet with our EMD²-based loss outperforms the original AlexNet slightly.

experiments.

The results are shown in Tab. 4. Our EMD²-based losses once again outperformed cross-entropy loss and L_2

loss (regression) significantly. We conducted additional experiments by discretizing the real-numbered aesthetic labels to 8 different number of bins (3,4,5,6,7,8,9,10 bins), which gave us 8 sets of ground truth labels. Then, we fine-tuned one VGG-EMD network for each set of ground truth and averaged the prediction results into an ensemble model, and denoted this method as VGG-EMD \times 8. It achieved state-of-the-art results trained only on image data, outperforming the previous state-of-the-art method trained with additional 11 labels such as color harmony and vivid color information.

4.5. Generalization on ImageNet

We show the generalization ability of using EMD² as a regularizer (Eq. 20) on the ImageNet ILSVRC 2012 dataset [31], which is a classification dataset with orderless-classes. We tested the original AlexNet [17] on this dataset with cross-entropy, and separately with the self-guided EMD² regularization. The results of the validations set are reported in Tab. 5. Although the AlexNet trained using our self-guided EMD² regularization was only slightly better, the results demonstrated the consistent improvements of our EMD²-based loss over the cross-entropy loss.

5. Conclusion

In this work, we argued that the conventional softmax cross-entropy loss for training DNNs under multi-class classification problems, only maximizes the predicted probability at the ground truth label, and ignores the inter-class relationships that exist in both ordered and orderless-classes datasets. We proposed to use the squared earth mover's distance (EMD²) as the loss function for the training to take class relationships into account, with an exact EMD² loss, and a self-guided EMD² regularization that automatically learns the ground distance matrix. We evaluated our methods on two age estimation datasets and one image aesthetic assessment dataset. Our method significantly outperformed state-of-the-art regression-based and cross-entropy-based CNNs using no external datasets, and with only image information. Furthermore, we generalized the EMD²-based loss to classification tasks with orderless-classes dataset, and showed performance improvement on ImageNet. Our future works include a more sophisticated ground distance matrix computing method, and exploring variants of EMDs that can perform better on orderless-classes datasets.

References

- [1] V. Belagiannis, C. Rupprecht, G. Carneiro, and N. Navab. Robust optimization for deep regression. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 2830–2838. IEEE, 2015. 1, 6
- [2] V. I. Bogachev and A. V. Kolesnikov. The monge-kantorovich problem: achievements, connections, and per-

- spectives. *Russian Mathematical Surveys*, 67(5):785, 2012. 2
- [3] J.-C. Chen, A. Kumar, R. Ranjan, V. M. Patel, A. Alavi, and R. Chellappa. A cascaded convolutional neural network for age estimation of unconstrained faces. In *BATS*, 2016. 6
 - [4] R. Datta, D. Joshi, J. Li, and J. Z. Wang. Studying aesthetics in photographic images using a computational approach. In *European Conference on Computer Vision*, pages 288–301. Springer, 2006. 7
 - [5] R. Datta, D. Joshi, J. Li, and J. Z. Wang. Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys (CSUR)*, 40(2):5, 2008. 7
 - [6] L. Deng and J. C. Platt. Ensemble deep learning for speech recognition. In *INTERSPEECH*, pages 1915–1919, 2014. 1
 - [7] C. Dong, C. C. Loy, K. He, and X. Tang. Learning a deep convolutional network for image super-resolution. In *European Conference on Computer Vision*, pages 184–199. Springer, 2014. 1
 - [8] Y. Dong, Y. Liu, and S. Lian. Automatic age estimation based on deep learning algorithm. *Neurocomputing*, 187:4–10, 2016. 1, 8
 - [9] E. Eidinger, R. Enbar, and T. Hassner. Age and gender estimation of unfiltered faces. *IEEE on IFS*, 2014. 1, 5, 6, 8
 - [10] S. Escalera, J. Fabian, P. Pardo, X. Baró, J. Gonzalez, H. J. Escalante, D. Misevic, U. Steiner, and I. Guyon. Chalearn looking at people 2015: Apparent age and cultural event recognition datasets and results. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 1–9, 2015. 1
 - [11] C. Frogner, C. Zhang, H. Mobahi, M. Araya, and T. A. Poggio. Learning with a wasserstein loss. In *Advances in Neural Information Processing Systems*, pages 2053–2061, 2015. 2
 - [12] A. C. Gallagher and T. Chen. Understanding images of groups of people. In *CVPR*, 2009. 1, 5, 7
 - [13] P. Golik, P. Doetsch, and H. Ney. Cross-entropy vs. squared error training: a theoretical and experimental comparison. 1, 3
 - [14] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates, et al. Deep speech: Scaling up end-to-end speech recognition. *arXiv*, 2014. 1
 - [15] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *CVPR*, 2015. 1
 - [16] S. Kong, X. Shen, Z. Lin, R. Mech, and C. Fowlkes. Photo aesthetics ranking network with attributes and content adaptation. *arXiv preprint arXiv:1606.01621*, 2016. 1, 5, 7, 8
 - [17] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 1, 5, 6, 8
 - [18] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 1
 - [19] G. Levi and T. Hassner. Age and gender classification using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 34–42, 2015. 6, 7
 - [20] E. Levina and P. Bickel. The earth mover’s distance is the mallows distance: some insights from statistics. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, volume 2, pages 251–256. IEEE, 2001. 3, 4
 - [21] H. Li, Z. Lin, X. Shen, J. Brandt, and G. Hua. A convolutional neural network cascade for face detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5325–5334, 2015. 1
 - [22] S. Liu, N. Yang, M. Li, and M. Zhou. A recursive recurrent neural network for statistical machine translation. In *ACL (1)*, pages 1491–1500, 2014. 1
 - [23] X. Lu, Z. Lin, H. Jin, J. Yang, and J. Z. Wang. Rapid: rating pictorial aesthetics using deep learning. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 457–466. ACM, 2014. 7
 - [24] D. G. Luenberger. *Introduction to linear and nonlinear programming*, volume 28. Addison-Wesley Reading, MA, 1973. 2, 4
 - [25] K. Ma, D. Samaras, M. Petrucci, D. L. Magnus, et al. Texture classification for rail surface condition evaluation. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9. IEEE, 2016. 1
 - [26] U. Park, Y. Tong, and A. K. Jain. Age-invariant face recognition. *IEEE transactions on pattern analysis and machine intelligence*, 32(5):947–954, 2010. 6
 - [27] N. Ramanathan and R. Chellappa. Face verification across age progression. *IEEE Transactions on Image Processing*, 15(11):3349–3361, 2006. 6
 - [28] K. Ricanek and T. Tesafaye. Morph: A longitudinal image database of normal adult age-progression. In *7th International Conference on Automatic Face and Gesture Recognition (FG06)*, pages 341–345. IEEE, 2006. 6
 - [29] R. Rothe, R. Timofte, and L. Van Gool. Deep expectation of real and apparent age from a single image without facial landmarks. *IJCV*, 2016. 1, 6, 7
 - [30] Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99–121, 2000. 2, 3
 - [31] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015. 2, 5, 8
 - [32] S. Shalev-Shwartz and A. Tewari. Stochastic methods for l1-regularized loss minimization. *Journal of Machine Learning Research*, 12(Jun):1865–1892, 2011. 2, 4
 - [33] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2014. 1, 5, 6
 - [34] R. Socher, C. C. Lin, C. Manning, and A. Y. Ng. Parsing natural scenes and natural language with recursive neural networks. In *ICML*, pages 129–136, 2011. 1
 - [35] D. Sutić, I. Brešković, R. Huić, and I. Jukić. Automatic evaluation of facial attractiveness. In *MIPRO, 2010 Proceedings of the 33rd International Convention*, pages 1339–1342. IEEE, 2010. 1

- [36] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, 2016. 6
- [37] D. Yu, K. Yao, H. Su, G. Li, and F. Seide. Kl-divergence regularized deep neural network adaptation for improved large vocabulary speech recognition. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 7893–7897. IEEE, 2013. 3
- [38] S. Zagoruyko and N. Komodakis. Wide residual networks. *BMVC*, 2016. 1, 5, 6, 7